

Is my Driver Observation Model Overconfident?

Input-guided Calibration Networks for Reliable and Interpretable Confidence Estimates

Alina Roitberg Kunyu Peng David Schneider
 Kailun Yang Marios Koulakis Manuel Martinez Rainer Stiefelhagen

Abstract—Driver observation models are rarely deployed under perfect conditions. In practice, illumination, camera placement and type differ from the ones present during training and unforeseen behaviours may occur at any time. While observing the human behind the steering wheel leads to more intuitive human-vehicle-interaction and safer driving, it requires recognition algorithms which do not only predict the correct driver state, but also determine their prediction quality through *realistic and interpretable confidence measures*. Reliable uncertainty estimates are crucial for building trust and are a serious obstacle for deploying activity recognition networks in real driving systems.

In this work, we for the first time examine how well the confidence values of modern driver observation models indeed match the probability of the correct outcome and show that raw neural network-based approaches tend to significantly overestimate their prediction quality. To correct this misalignment between the confidence values and the actual uncertainty, we consider two strategies. First, we enhance two activity recognition models often used for driver observation with temperature scaling – an off-the-shelf method for confidence calibration in image classification. Then, we introduce Calibrated Action Recognition with Input Guidance (CARING) – a novel approach leveraging an additional neural network to learn scaling the confidences *depending on the video representation*. Extensive experiments on the Drive&Act dataset demonstrate that both strategies drastically improve the quality of model confidences, while our CARING model outperforms both, the original architectures and their temperature scaling enhancement, leading to best uncertainty estimates.

Index Terms—Driver activity recognition, model confidence reliability, uncertainty in deep learning.

I. INTRODUCTION

WITH the rapidly growing accuracy of driver observation models [1], the existing gap between the published methods and their applications in practice makes us wonder about important performance aspects potentially being overlooked. When examining the previous research, we make two observations. First, the existing driver activity recognition algorithms are highly driven by the top-1 accuracy [1]–[8], skipping other relevant metrics, such as the *reliability of their confidence values*. Second, recent fundamental deep learning

This work was supported by the Competence Center Karlsruhe for AI Systems Engineering (CC-KING, www.ai-engineering.eu) sponsored by the Ministry of Economic Affairs, Labour and Housing Baden-Württemberg.

Authors are with Institute for Anthropomatics and Robotics, Karlsruhe Institute of Technology, Germany (e-mail: firstname.lastname@kit.edu).

Corresponding author: Alina Roitberg.

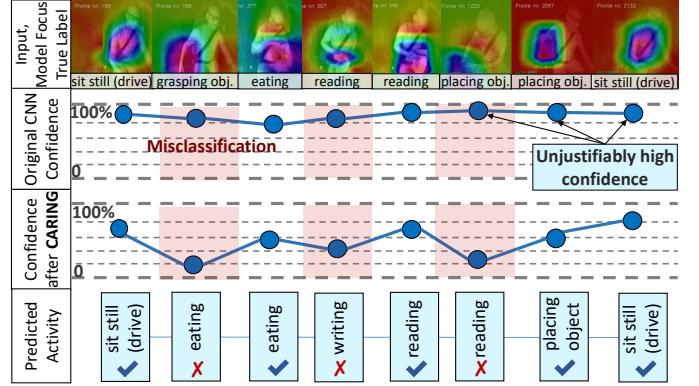


Fig. 1: Deep neural networks have seen a surge in popularity in driver observation, but worrying evidence indicates that such models produce unjustifiably high-confidence predictions. Overconfident models imply that if a network predicts certain secondary behaviour with $> 99\%$ confidence (which is indeed very common), the prediction is oftentimes correct in much less than in 99% of cases. This figure depicts an example of observed predictions and confidence values of a driver activity recognition CNN during a single driving session. If we directly use the probability estimates provided by the original CNNs, we end up with *unjustifiably high confidences*, which can be very damaging in practice. Temperature-based methods [9] apply a fixed scaling to the CNN probability estimates in order to obtain more realistic values. Our approach, CARING, goes one step further, and dynamically infers a suitable scaling factor depending on the current observation, leading to more reliable confidences.

research raises alarm about neural networks being notably bad at detecting ambiguities [9]–[12]. Despite the impressive gain in accuracy on modern driver activity recognition datasets [1], [5], [13], examining how well the confidence values of driver behaviour understanding models indeed reflect the probability of a correct prediction (see Figure 1) has been overlooked in the past and is the main motivation of our work.

Humans have a natural intuition about probabilities [14]: If a neural network reports that the driver is “talking on the phone” with 99% confidence, we tend to take this prediction for granted. In neural networks, the output of the last fully-connected is usually normalized through the *Softmax*¹ function

¹Softmax function is often applied on *logits* (the output vector of the last layer of a classification network, where each value represents a category score) and normalizes them to sum up to one by computing the exponents of each output and then normalizing each of them by the sum of those exponents.

to sum up to one, therefore *resembling* prediction probabilities. However, blindly interpreting these values as probabilities of the outcome being correct would be naive, as the Softmax scores are not calibrated and merely optimized for a high top-1 accuracy on a static set driver of behaviours. The output Softmax vector *appears* to hold probabilities of the individual classes, but they do not necessarily match the true confidence of the model [9], [15]. Furthermore, worrying evidence from deep learning research highlights the problem of *model miscalibration*, suggesting that raw neural network confidences tend to be biased towards very high values for both, correct and incorrect predictions [9], [15]. In real-life driving applications, on the other hand, models need to be more reserved: wrong high-confidence decision can lead to tragic outcomes. At the same time, indicating high classification uncertainty and reacting accordingly is the better choice. In this work, we equip driver observation frameworks with building blocks for reliably quantifying model confidence.

We aim to highlight the need for interpretable and reliable uncertainty measures as a means for identifying misclassifications inside the vehicle cabin. Although this area is seeing growing attention in the field of general image recognition [9], [11], it is underresearched in video classification and especially driver observation. To achieve this, we integrate the *reliability* of model confidence in the large-scale driver activity recognition benchmark *Drive&Act* [1] and develop models which learn to convert oftentimes overconfident uncertainty values of the original networks into reliable probability estimates.

Why driver activity recognition? Rising levels of automation increase human freedom, leading to drivers being engaged in distractive behaviors more often while the type of activities become increasingly diverse. For example, *working on laptop* or *reading magazine* behind the steering were almost unthinkable until now, but these behaviours become more common as the driver is gradually relieved from actively steering the car. Although distractions become safer as the vehicle becomes more intelligent, this change does not happen from one day to another and is a rather long-lasting transformation [16]. Over-reliance on artificial intelligence might lead to catastrophic consequences, and, for a long time, the driver will need to intervene in case of uncertainty [16]–[18]. However, there are important long-term application scenarios of driver activity recognition even in fully-autonomous cars. For example, understanding the situation inside the vehicle cabin may increase comfort, e.g., by adjusting the driving style if the person is drinking a hot drink or being an intuitive communication interface via gestures. Applications of driver activity recognition models depend on the degree of vehicle automation [19].

We summarize four main use-cases for applications of driver activity recognition [1], [20], [21]:

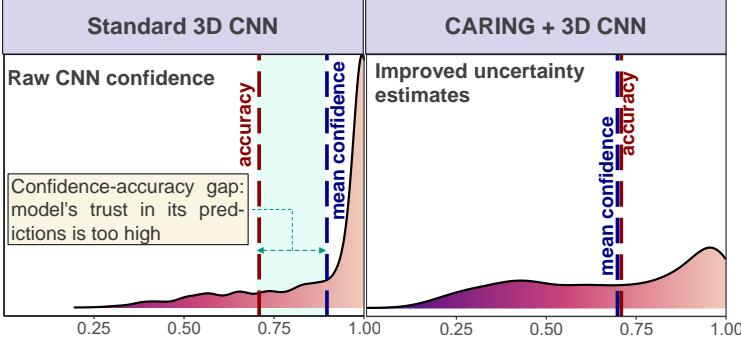
- **Improved safety through identified distraction.** Recent studies highlight that current activity directly affects human cognitive workload in both, general- and driving context [22]–[24]. For example, Deo and Trivedi [24] suggest that certain secondary activities such as *interacting with the infotainment unit* negatively impact the perceived

readiness-to-take-over. Therefore, the key application of such algorithms at SAE levels 0 to 3 [19] is the assessment of human distraction and reacting accordingly, for example, with a warning signal.

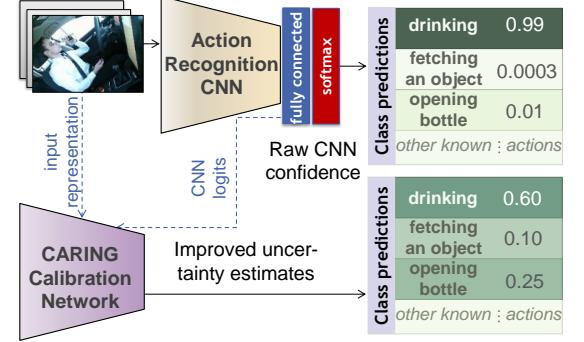
- **Increased comfort through automatic driver-centered adaptation.** With the automation rising to SAE levels 4 and 5, increasing driver comfort by automatic adaptation of the vehicle controls becomes the more important use-case. For example, movement dynamics might automatically adjust depending on the detected activity, e.g., softer driving if the person is drinking tea or sleeping.
- **Novel intuitive communication interfaces.** Visually recognizing driver gestures might lead to novel communication interfaces inside the vehicle, serving as a more intuitive alternative for the central console, as previous research identifies hand signals as a highly convenient way for human-machine interaction [2], [25], [26].
- **Identifying and preventing dangerous maneuvers.** Since the majority of traffic fatalities are caused by inappropriate driving events induced by humans [27], [28], a further safety-related application of driver activity recognition during manual driving is intention prediction. Timely anticipation of driver intention offers a possible solution to prevent potential accidents at an early stage, allowing ADAS to notice that the person e.g. is going to induce a dangerous turn and prevent the accident by taking over the control or notifying the driver [29].

A Note on Terminology While various meanings have been given to the term *uncertainty*, we specifically refer to the “classification uncertainty of a discriminative classification model” as to *inverse of model’s confidence* [30], [31]. We follow the probabilistic definition of *model’s confidence* [9], [32]–[34], which is a numeric value depicting a probability estimate of a class membership in the context of multi-class categorization. Lastly, the *reliability of model confidence* represents the quality of the observed confidence estimates, where we adopt the mathematical formulation of [9], which intuitively means that if a network produces a confidence value of 95% it should indeed be correct 95% of times, (i.e. this would also be the value of model’s accuracy). The exact mathematical formulation of this definition is provided in Section III-A.

Contributions and Summary This paper makes the first step towards driver observation models capable of identifying misclassifications by realistically estimating their prediction confidence and has following major contributions. (1) We incorporate the Expected Calibration Error and the Brier Score metrics in the evaluation protocol of the large scale driver activity recognition dataset *Drive&Act* [1] and benchmark two CNN-based architectures commonly-used for driver observation in terms of their confidence reliability. Our experiments indicate, that raw driver activity recognition models do not provide reliable uncertainty estimates and *strongly overestimate their confidence*. To the best of our knowledge, this is the first study of the relation between the confidence of a driver activity recognition model and the likelihood of a correct prediction. (2) We augment driver activity recognition



(a) An empirical example of experienced confidence-accuracy gap of modern video classification CNNs used for driver activity recognition. We visualize the distribution of the Softmax confidence values of the popular Pseudo 3D ResNet model deployed on the Drive&Act dataset (validation split). The X-axis represents the space of possible probability estimates, while the Y-axis illustrates the frequency of samples with the corresponding probability estimate before (left) and after (right) the improvement through our model. Dashed red and blue lines depict the average confidence values and the accuracy accordingly which would match in a model with perfectly calibrated confidences. Native confidence values clearly underestimate model uncertainty (most samples received a confidence value of > 90%, while the accuracy is much lower). We therefore propose to incorporate the **reliability** of model confidences in the evaluation of driver observation models and introduce algorithms for improving it.



(b) An overview of the proposed Calibrated Action Recognition with **Input Guidance** (CARING) model for obtaining more realistic uncertainty estimates. CARING leverages an additional calibration neural network for learning to obtain realistic confidence estimates *depending on the input*. The additional calibration network is optimized on a held-out validation set, during which the main classifier is frozen. It takes as input a video representation (intermediate features of the main classifier) and learns to infer an instance-dependent scaling factor, which is then used to scale the logits of the original network, therefore "softening" the resulting confidences and leading to much more realistic confidence estimates for driver observation.

Fig. 2: Observed problem of overconfident driver observation CNNs and an overview the proposed model for mitigating this effect.

models with the temperature scaling approach [9] – an off-the-shelf approach for model confidence calibration in image classification. This approach learns *one global scalar* τ through which the network logits are divided to obtain better confidence estimates. (3) We expect, that taking the video input representation into account gives us cues which are helpful for estimating prediction uncertainty, and present a new method named **Calibrated Action Recognition with Input Guidance** (CARING). A novel aspect of our approach is that CARING entails an additional calibration network, which learns to produce flexible temperature values $T(z)$ *dependent on the input z*. In contrast, [9] uses one global scalar for scaling the confidences, so that the logits are always divided by the same value, while we dynamically infer the temperature values which suit currently observed state.

Despite disregarding the input, temperature scaling alone leads to much more meaningful confidence values, while our input-dependent CARING model excels these improvements, reducing the expected calibration error to only $\sim 5\%$ on Drive&Act and to only $\sim 2.5\%$ in case of sufficient training data. However, our results also showcase room for improvement in case of underrepresented categories where the confidence-accuracy mismatch remains high even after significant improvement through CARING.

This paper is an extension of our conference publication [35], which has been expanded with specific focus on *driver* behavior understanding, a detailed description of the proposed CARING method, and an extended set of experiments and analyses.

II. RELATED WORK

A. Driver Action Recognition

Driver activity recognition is an essential task in Advanced Driving Assistance Systems (ADAS), aiming at recognition

of normal- and abnormal driver behaviors, such as eating, phoning, and smoking. According to [36], over 90% of vehicle crashes are due to drivers' abnormal behaviors and driving errors in the US, while the total number of traffic accidents might be decreased by 10 – 20% if a good driving assistant monitoring system would be in place [37]. Similar to other computer vision fields, existing research works in driver action recognition field can be also divided into two main groups: (1) approaches based on manually defined features and (2) deep learning based approaches, which learn intermediate representations end-to-end.

Feature-based approaches. Until recently, the majority of driving observation frameworks comprised a manual feature extraction step followed by a classification module (for a thorough overview see [21]). The constructed feature vectors are often derived from hand- and body-pose [2], [3], [6], [7], [38], [39], facial expressions and eye-based input [40], [41], and head pose [42], [43], but also foot dynamics [44], detected objects [6], and physiological signals [45] have been considered. Classification approaches are fairly similar to the ones used in standard video classification, with LSTMs [3], [4], SVMs [2], [46], random forests [47] or HMMs [4], and graph neural networks [7], [48] being popular choices.

Deep learning-based approaches. Since the rise of deep learning, approaches based on Convolutional Neural Networks (CNNs), which operate directly on the image input and learn intermediate representations end-to-end, took over the top of most computer vision benchmarks, ranging from object classification and -detection to action recognition and semantic segmentation [49]–[55]. These developments also had strong influences on the field of driver observation where a variety of works report top performance of CNN-based architectures [1],

[5], [13], [37], [56]–[62], with spatiotemporal CNNs, such as I3D [52], P3D [51], and 3D-MobileNet [63] being popular backbones [1], [13], [60], [61], [64]. Further accuracy improvements are often achieved through techniques such as computation and fusion of additional modalities (e.g., semantic segmentation, optical flow) [8], [37] or use of attention modules [5], [36] and CNN-RNN combinations [8], [65], [66].

The general goal of the above research is to improve the recognition accuracy on driver activity recognition datasets without considering whether the network *confidence indeed reflects the likelihood of a correct prediction* (Figure 4 demonstrates the confidence-accuracy mismatch of P3D on the Drive&Act [1] dataset revealed by our study).

In our work, we go beyond the incentive of high top-1 accuracy and aim for driver activity recognition models with *realistic confidence estimates*. Keep in mind, that while our method leads to far more realistic uncertainty estimates, the accuracy is not impacted, as the transformation scales the output without changing the category order.

B. Detecting Model Misclassifications

Since real applications of deep CNNs need good quality of confidence measures, several researchers have highlighted this requirement [10]–[12]. However, this issue still lacks attention in both, driver- and general activity recognition. Still, multiple works have addressed this topic in classical machine learning [67]–[69], image classification [9], [11], [15], and person identification [70], offering useful insights for applying these techniques for video-based driver observation. A group of those methods follows the Bayesian approach, such as Monte Carlo Dropout sampling [15] or using an ensemble of models to estimate uncertainty [71]. Compared to dropout sampling, which gives gaussian distributions to assess model uncertainty, calibration-based techniques [9], [69], [72]–[75] return a single confidence value. Calibration-based methods hold several advantages: they do not require sampling and are therefore less computationally expensive and their result is *calibrated*, meaning that the resulting values ideally match human interpretation of probabilities (i.e. the likelihood of a correct outcome). Guo et al. [9] compared several calibration-based approaches such as isotonic regression, histogram binning, and Bayesian quantile binning regression on image and document classification tasks. In their work, they introduced *temperature scaling* – an approach, which is a variant of the Platt Scaling approach [69], in which a single global parameter is learned on a validation set, and then utilized to scale the network logits and make the output confidence closer to the actual accuracy. Temperature scaling was simple but proved to be highly effective, outperforming far heavier approaches [9] and later found success in natural language processing [74] and medical applications [76].

We for the first time integrate different building blocks for identifying the faithfulness of model confidences for driver activity analysis. In particular, we augment driver activity recognition networks with uncertainty estimation through (1) the raw Softmax confidences [11], (2) the temperature scaling technique of Guo et al. [9], and (3) our newly proposed CARING model, which, in contrast to [9], features *input-driven*

confidence transformation. While in temperature scaling, the logits are always divided by the same global value, our CARING model dynamically infers different scaling parameters for different inputs through an additional calibration network.

III. UNCERTAINTY-SENSITIVE ACTION RECOGNITION

We aim to explore *reliability of model confidence* in the field of driver observation, where the existing frameworks have been mostly motivated by high accuracy [1], [2], [37], [62]. First, we formalize the problem and equip the *Drive&Act* benchmark [1] with multiple metrics linked to the confidence quality (Section III-A). We then combine two 3D CNNs (model description in Section III-B) often used for driver activity recognition, with three different strategies for quantifying prediction confidence: using raw *Softmax* output (Section III-C), transforming the confidences through temperature scaling (Section III-D) and our newly proposed CARING model, which learns to infer optimal scaling parameters depending on the input (Section III-E).

A. Problem Definition: Reliable Confidence Measures

Intuitively, the confidence of a classification model should reflect its accuracy, i.e., *reliable* confidence should indicate the probability of a prediction to be correct. Based on a driver observation x , we predict an action class a_{pred} with model confidence $\hat{p}(a_{pred}) \in [0, 1]$, $a \in A_{\{1, \dots, m\}}$ being the set of possible driver behaviour classes. Provided with the ground-truth label a_{true} , we consider our model confidence to be well calibrated and reliable if it, on average, reflects the probability of a successful outcome $\mathbb{P}(a_{pred} = a_{true})$. We follow the definition of [9] and assume that the confidence quality of a model is perfect if the following condition is met:

$$\mathbb{P}(a_{pred} = a_{true} | \hat{p}(a_{pred}) = p) = p, \quad \forall p \in [0, 1] \quad (1)$$

The model therefore delivers reliable probability estimates if the mismatch between the average model's confidence and the prediction accuracy is low (the right side of Figure 2 illustrates such confidence-accuracy disarray).

In practice, perfect model calibration and even its *perfect* evaluation are not possible, since Equation (1) is defined on the continuous interval $[0, 1]$ while we only have a finite amount of confidence estimates $\hat{p}(a_{pred})$. However, this value can be *approximated* with different techniques, e.g., discrete partitioning of the probability space. We consider three different metrics: Expected Calibration Error (ECE) [9], the presumably most popular metric for quantifying confidence discrepancy, the normalized Brier score [77], and Negative Log Likelihood.

Expected Calibration Error (ECE). One way to solve the issue of continuous probability space in Equation (1) is quantizing the $[0, 1]$ interval into equally sized segments $seg_i = [\frac{i}{K}, \frac{i+1}{K})$ with $i \in 0, \dots, K - 1$ ($K = 10$ in this work) and calculating the average accuracy and confidence among all N_{seg_i} samples falling into the same confidence segment (see Figure 3). In case of perfect confidence calibration, the accuracy-confidence distance of the individual segments should be zero. The Expected Calibration Error (ECE) [9],

leverages such partitioning and measures the absolute difference between accuracy $acc(seg_i)$ and average confidence $\hat{p}(seg_i)$ for each individual segment seg_i , followed by averaging among all segments weighted by the number of samples in each of the K segments. Formally, ECE is computed as:

$$ECE = \sum_{i=1}^K \frac{N_{seg_i}}{N_{total}} |acc(seg_i) - \hat{p}(seg_i)|, \quad (2)$$

where N_{total} denotes the number of samples in all segments. In deep learning research, ECE is perhaps the most widespread measure of model confidence reliability [9], [75], [76], [78], and we therefore also view it as our primary performance metric.

Reliability Diagrams. ECE can be effectively visualized using *reliability diagrams* (Figure 3), where the X-axis represents the probability space discretized into K bins. In each bin, we visualize a bar with size matching the empirically computed accuracy of samples, which confidence fell into this segment. Since the accuracy and the confidence match in perfect case, the bars should ideally match the diagonal. Our results in Figure 3 showcase, that for an out-of-the-box (illustrated on the left), the accuracies are significantly *below* the confidence value. The original model is therefore *overly confident*.

Brier Score. Another metric for quantifying model confidence quality is the Brier score [77], initially developed to quantify weather forecasts reliability and later adapted as a general confidence reliability measure. The Brier score represents squared error between the estimated probabilities and the one-hot encoding of the ground-truth data. As the original definition has different boundaries for binary and multi-class cases [77], we use the normalized version of the Brier score, where the values lie between 0 and 1, formally defined as [79]:

$$Brier = \frac{1}{2N_{total}} \sum_{i=1}^{N_{total}} \sum_{a=1}^m (C_{a,i} - \hat{p}(a,i))^2, \quad (3)$$

where N_{total} is the total number of samples, m is the number of possible driver activity classes, $\hat{p}(a,i)$ is the model probability estimates of the i th sample belonging to category a and $C_{a,i} \in \{0, 1\}$ is 1 if a is the ground-truth class of the i th sample, otherwise 0. Note, that previously published research describes Brier score as a metric is *insensitive towards rare classes* in the test set [78] and should be taken with caution since in real-life datasets the categories are rarely equally distributed (which is also the case in *Drive&Act*).

Negative Log Likelihood (NLL). Negative Log Likelihood, which is also leveraged for optimizing the classifier itself, can be used to measure quality of the uncertainty estimates [9] and is therefore reported in our study.

B. Backbone Neural Architectures

We evaluate model confidence reliability on two widely used architectures, the Pseudo 3D ResNet [51] as well as the Inflated 3D ConvNet [52]. Both networks have proven to be successful in driver observation [1], [13]. While the P3D

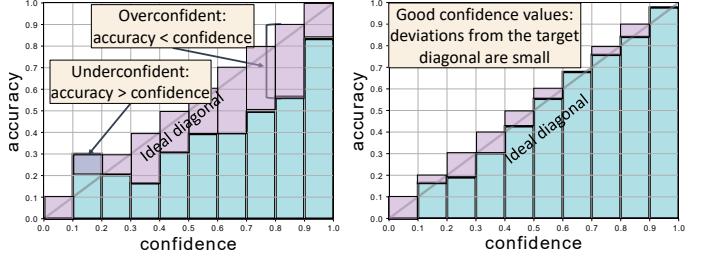


Fig. 3: Reliability diagrams of a model with poor confidence estimates (left) and a well-calibrated model (right). The illustrated data are the confidence values of the Pseudo 3D ResNet at the Drive&Act validation split before and after the improvement with the CARING calibration network.

ResNet decouples spatial and temporal convolutions by either using a $(3 \times 3 \times 1)$ filter kernel for the spatial domain or a $(1 \times 1 \times 3)$ filter kernel for the temporal domain, the I3D Network makes use of filter kernels which perform spatial and temporal convolutions at the same time. Another structural difference, is that P3D ResNet has an impressive depth of 152 layers enabled through use of residual connections to ease the gradient flow, while I3D has 27 layers and stacks multiple Inception modules, which perform different types of convolutions in parallel.

C. Softmax Output as Confidence Measures

Similar to other neural networks, the last fully-connected layer are referred to as a *logit vector* \mathbf{y} . The individual values y_a are scores marking how likely the activity a is indeed the current class. The raw *logit vector* is *not normalized*. In order to mimic a probability function (have values between 0 and 1 and sum up to 1), it is commonly normalized using *Softmax*. Although these scores are rather *pseudo-probabilities*, they are often used as model uncertainty estimates [11]. The confidence estimates in this case are obtained as:

$$\hat{p}(a_{pred}) = \max_{a \in \mathcal{A}} \frac{\exp(y_a)}{\sum_{\tilde{a} \in \mathcal{A}} \exp(y_{\tilde{a}})} \quad (4)$$

D. Calibration via Temperature Scaling

In image recognition, an off-the-shelf strategy for improving confidence estimates is *temperature scaling* [9]. Temperature scaling leverages a single, global parameter τ to scale the logit vector prior to *Softmax* normalization, with the final estimates becoming:

$$\hat{p}(a_{pred}) = \max_{a \in \mathcal{A}} \frac{\exp(y_a / \tau)}{\sum_{\tilde{a} \in \mathcal{A}} \exp(y_{\tilde{a}} / \tau)} \quad (5)$$

This introduces the properties $\lim_{\tau \rightarrow \infty} \hat{p}(a_{pred}) = 1/m$ and $\lim_{\tau \rightarrow 0} \hat{p}(a_{pred}) = 1$, meaning that τ is a parameter which either decreases confidence ($\tau > 1$) or increases confidence ($\tau < 1$). After the classifier is trained to assign the categories in a standard way, the model weights are frozen and the parameter τ is optimized on a held-out validation set also using Negative-Log-Likelihood. Note, that this scaling affects

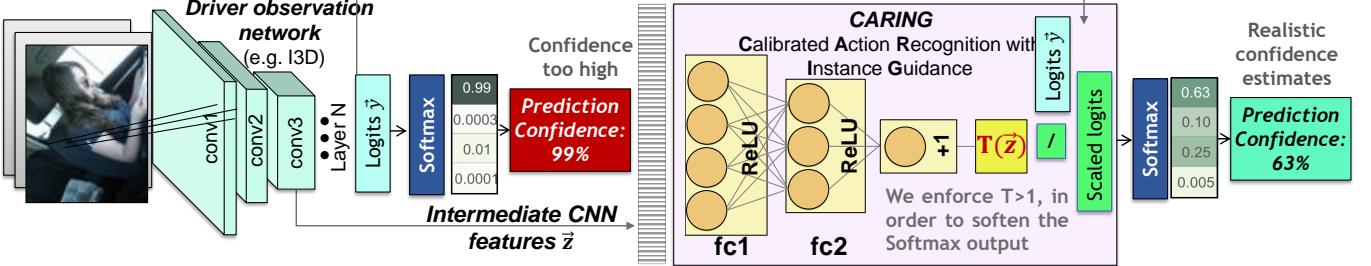


Fig. 4: Overview of the Calibrated Action Recognition under Instance Guidance Model (CARING). CARING is an additional neural network which learns to infer the scaling factor T depending on the instance representation. The logits of the original activity recognition network are then divided by T , giving better estimates of the model uncertainty. In our implementation, the intermediate vector \mathbf{z} is extracted from the 9th inception block of the I3D model, while for P3D we use the layer before the last fully-connected layer.

probability estimates of all classes equally and does not alter the order of the top-k categories or the final top-1 prediction of the model. This means that the accuracy of the model stays the same. Although the method is remarkably simple, it had been extremely effective for obtaining well-calibrated image recognition CNNs, surpassing more complex approaches [9]. We learn a suitable temperature parameter τ using Gradient Descent with a learning rate of 0.01 for 50 epochs.

E. Calibrated Action Recognition with Input Guidance

Next, we introduce a new approach for obtaining good uncertainty estimates by learning to scale the logits *depending on the input*. While in the next section our experiments will reveal, that temperature scaling significantly improves the confidence reliability on its own, it *does not consider the current observation* when transforming the confidences, as the logits are always divided by the same temperature scalar τ . Our idea is leveraging information present in the current driver observation, as it may carry important hints about the prediction reliability (i.e. the sample might contain visible noise). Therefore, we propose to infer an individual input-dependent scaling factor $\tau(\mathbf{z})$ during inference instead of learning a single global parameter τ , resulting in the confidence calculation. Our model and the temperature scaling approach [9] therefore have one crucial difference: the scaling factor is not fixed but varying dependent on current input. We therefore dynamically infer the scaling parameter $\tau(\mathbf{z})$ on-the-fly at test-time depending on the video representation \mathbf{z} . Therefore, the transformed logits become $\mathbf{y}_{scaled} = \mathbf{y}/\tau(\mathbf{z})$.

To enable such input-dependent inference of $\tau(\mathbf{z})$, we propose an additional *calibration neural network*, which we refer to as **CARING** (**C**alibrated **A**ction **R**ecognition under **I**nput **G**uidance), as it guides the logit transformation depending on the video input. An overview of our model is provided in Figure 4. CARING is a lightweight regression network with two fully-connected layers and the output of the second layer being a single regressed value. The final scaling factor is obtained by applying the *relu* activation and adding 1 to assure $\tau(\mathbf{z}) \geq 1$, so that the probability estimates are softened. The input-dependent temperature $\tau(\mathbf{z})$ is therefore obtained as:

$$\tau(\mathbf{z}) = 1 + \text{relu}(\mathbf{W}_2 \text{ relu}(\mathbf{W}_1 \mathbf{z} + \mathbf{b}_1) + \mathbf{b}_2), \quad (6)$$

where $\mathbf{W}_1, \mathbf{W}_2$ depict the weight matrices, and $\mathbf{b}_1, \mathbf{b}_2$ are the bias terms, and \mathbf{z} is an intermediate representation of current

input. We leverage the intermediate feature vector of the backbone classification network, which have dimensionality of 1024 and 2048 for Inflated 3D ConvNet and Pseudo 3D ResNet, respectively.

After we regress the input-dependent temperature $\tau(\mathbf{z})$, the logits are divided by it before the *Softmax* function and our final confidence estimates become:

$$\hat{p}(a_{pred}) = \max_{a \in \mathcal{A}} \frac{\exp(\frac{y_a}{\tau(\mathbf{z})})}{\sum_{\tilde{a} \in \mathcal{A}} \exp(\frac{y_{\tilde{a}}}{\tau(\mathbf{z})})}. \quad (7)$$

Training procedure. Conceptually, our proposed architecture comprises two main components: 1) the standard classification network aimed at assigning the correct driver behaviour to the input video and 2) an additional calibration network used to scale the classifier logits and produce more realistic confidences. First, we train the classifier following the standard procedure and Negative-Log-Likelihood (NLL) optimization. In the second step, the classifier weights are frozen and the calibration network is optimized on a separate, held-out validation set which examples were not present during the classifier training. Different losses can be used for training such network, including the NLL loss [9], Maximum Mean Calibration Error (MMCE) [80], Brier Loss [77] and Focal Loss [81]. We selected the NLL loss in our training, since it is presumably the most popular choice [9], [74], [76] for confidence calibration and is also used in the temperature scaling approach of [9] to which our model is compared.

By freezing the model weights and only training our lightweight calibration model on a validation set, a realistic distribution of correct and incorrect predictions is provided and being overly confident about these predictions would increase the loss. In the same fashion as temperature scaling, CARING is a post-processing technique for improving model probability estimates. It *does not influence the accuracy*, as the order of the logit categories does not change (we do divide by different temperature values for different observations, but all the logit vector values are divided by the same scalar). We optimize CARING with a learning rate of 5e-3 and a weight decay of 1e-6 for 300 epochs on a held-out validation set using the previously mentioned NLL loss while maintaining the weights of the backbone classification network frozen. The complete process is described as pseudo-code in Algorithm 1.

```

Training of I3D
train, valid, test = split(Drive & Act)

for epoch in [1...300]:
    for data, labels in train:
        logits = i3d(data)
        loss = NLL(logits, labels)
        backpropagate(loss)

Training of CARING
caring = Sequential(
    Linear, ReLU,
    Linear, ReLU,
    Lambda(x -> x + 1))

for epoch in [1...300]:
    for data, labels in valid:
        logits = i3d(data)
        intermediate_features = truncate(i3d)(data)
        temperature = caring(intermediate_features)
        scaled_logits = logits / temperature
        loss = NLL(scaled_logits, labels)
        backpropagate(loss)

Inference and evaluation with ECE
for data, labels in test:
    logits = i3d(data)
    intermediate_features = truncate(i3d)(data)
    temperature = caring(intermediate_features)
    scaled_logits = logits / temperature
    logits_i3d.append(logits)
    logits_caring.append(scaled_logits)
    labels_all.append(labels)

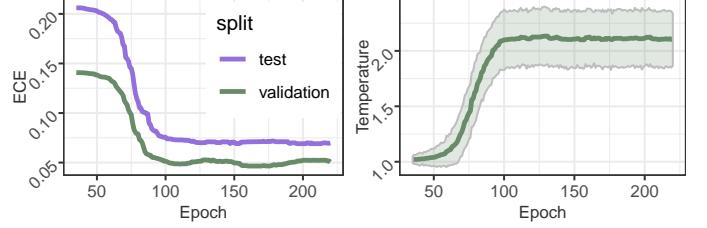
(See Eq. 2 for ECE computation)
ECE_i3d = ECE(logits_i3d, labels_all)
ECE_caring = ECE(logits_caring, labels_all)

```

Algorithm 1: The CARING training and evaluation procedure in the specific case of using I3D as a backbone and evaluation using the Expected Calibration Error (ECE) metric.

For the intermediate representation vector \mathbf{z} , we leverage the output of the 9th inception block of I3D. We average the output over the temporal dimension to reduce the vector size, resulting in 1024-dimensional vectors. For P3D, we use the layer before the last fully-connected layer (the output size is 2048). Furthermore, we observe that the CARING performance does not alternate significantly if the representations were extracted at medium or late network depths, but the quality declines at earlier layers. Note, that the size of \mathbf{z} directly impacts computational complexity of CARING. For our configuration, (*i.e.*, 1024-dimensional \mathbf{z} in case of I3D), CARING comprises 66.7 parameters and requires 0.13 MFLOPS for a forward pass, which is a small overhead given the significantly larger size of the classification network (12.06 M parameters and 107.9 GFLOPS for I3D).

Why is training the calibration network with NLL loss effective against miscalibration? Since both the classifier backbone and the CARING model use the NLL loss, we need to clarify why such optimization leads to unrealistic confidences after the standard classifier training but gives proper confidence estimates after the calibration. NLL itself indeed reflects confidence miscalibration [9] and we therefore use it as one of our metrics. The main reason for the confidence-accuracy disarray rising during the classifier training is *overfitting on the training*



(a) Expected Calibration Error improvement during the training procedure for standard deviation estimated by validation and test data.

Fig. 5: CARING model evolution during training for one Drive&Act split. Both average value and standard deviation of the learned input-dependent scaling parameter $\tau(z)$ rise as the training proceeds (right figure). Jointly with the decrease of the calibration error (left figure), this indicates the usefulness of learning different scaling parameters for different inputs.

set combined with the target labels being exclusively 0 or 1 [9], [80]. Target labels for classification training are commonly 0 or 1 (one-hot encodings of the ground truth). The high capacity of modern neural architectures leads to overfitting to the one-hot encoded labels, which is considered the key reason for model miscalibration [9], [81] (although other factors have been witnessed, such as batch normalization). By training a model to learn *Softmax* scores which mimic binary targets, it is incited to provide overly confident results. After the classifier training has converged in terms of *accuracy* (which is usually very high on the training data), it further optimizes the NLL-criteria to match the 0/1 labels, which will usually be 1 for the predicted class on the training data, leading to overly confident models. Consequently, the classifier becomes overly confident due to the unrealistically high amount of the correct 1-label predictions. By freezing the model weights and only training our lightweight CARING model on a held-out set, we bypass the issue of overfitting to 0/1 labels due to high model capacity. A realistic distribution of correct and incorrect predictions is therefore provided and being overly confident about these predictions would increase the loss. Note that overfitting is prevented since CARING does not influence the model accuracy (since the classifier is frozen and the logits order stays the same).

Does our CARING model indeed scale different inputs differently? Although CARING always takes in representation of the video input, it could theoretically ignore it and fall back learning a constant value, therefore converging to conventional temperature scaling. In order to validate if our model actually infers different $\tau(z)$ for different inputs, we analyze the Expected Calibration Error (Figure 5a) in combination with the mean and standard deviation of the scaling factor $\tau(z)$ (Figure 5b) computed over all validation samples during each CARING training epoch. Figure 5 reveals that while the ECE decreases during optimization, the standard deviation of the individual scaling factors increases, indicating that varying $\tau(z)$ is indeed beneficial for confidence reliability, which will also be validated empirically in the upcoming section.

IV. EXPERIMENTS

In this section, we for the first time benchmark modern driver activity recognition models with regard to their ability to identify failure cases through their own probability estimates. To this intent we for the first time consider metrics described in Section in the context of driver activity recognition [III-A](#). In a variety of settings featuring different input modalities and amount of training data, we benchmark two prominent video classification CNNs, against their versions equipped with the temperature scaling method [\[9\]](#) and our proposed CARING approach. We hope that by establishing this benchmark we will motivate researchers to rethink the roles of uncertainty estimates in driver observation.

This Section is structured as follows. First, we give an overview of the evaluation testbed and the underlying data (Section [IV-A](#)). Next, we provide a thorough evaluation of confidence estimates for driver activity recognition, focusing on the main evaluation setting of Drive&Act (a single NIR modality) and distinguishing between different amount of available training data (Section [IV-B](#)). Since different activity types might be relevant for different driving applications, Section [IV-C](#) is devoted to individual category-specific outcomes. Then, in Section [IV-D](#) we focus on recognition from different sensor types and placements within the vehicle cabin, demonstrating generalization capabilities of our model. Finally, Section [IV-E](#) provides qualitative analyses of the resulting distributions of model confidences.

A. Benchmark settings

Currently, there is no established evaluation scheme targeting the reliability of confidence values in the context of driver observation. Thereby, we adapt the evaluation protocols of Drive&Act [\[1\]](#) – a large-scale driver activity understanding testbed, to our task. Drive&Act [\[1\]](#), is a multimodal benchmark for driver activity recognition covering nine synchronized views and three modalities (near-infrared, depth, and color). Drive&Act comprises 34 fine-grained activity classes on the main evaluation level. Yet, the categories are heavily unbalanced, which is typical for application-specific datasets: the number of samples ranges from only 19 examples of *taking laptop from backpack* to 2797 instances of *sitting still*. Considering that CNNs often suffer when learning from few examples, the behaviors are sorted according to their frequency in the dataset and separated into *common* (top half of the classes) and *rare* (the bottom half) ones. The models are subsequently evaluated in three different modes: incorporating *all activities* as in the usual case, using only the *overrepresented*-, or only the *rare* categories.

For P3D- and I3D models, the input are snippets of 64 consecutive frames. If the original video segment is longer, the snippet is randomly chosen during training and at the center of the video at test-time. If the segment is shorter, the last frame will be repeated until the 64-frame snippet has been filled.

Following the problem formalization in Section [III-A](#), the standard accuracy-driven evaluation protocols [\[1\]](#) have been extended with the expected calibration error (ECE), which depicts the divergence of model confidence score from the

Model	ECE (%)		Brier		NLL	
	val	test	val	test	val	test
Drive&Act - Common Classes						
P3D [51]	16.9	19.39	0.27	0.32	1.63	1.85
I3D [52]	10.22	13.38	0.12	0.16	0.90	1.27
P3D + Temperature Scaling [9]	5.65	5.70	0.25	0.29	1.28	1.48
I3D + Temperature Scaling [9]	5.31	6.99	0.11	0.15	0.57	0.83
CARING - P3D (ours)	4.81	4.27	0.24	0.28	1.19	1.42
CARING - I3D (ours)	2.57	5.26	0.10	0.15	0.50	0.78
Drive&Act - Rare Classes						
P3D [51]	31.49	37.25	0.39	0.50	3.43	4.68
I3D [52]	31.48	43.32	0.34	0.47	3.41	4.54
P3D + Temperature Scaling [9]	17.83	21.09	0.34	0.42	2.26	2.99
I3D + Temperature Scaling [9]	24.97	32.38	0.30	0.41	1.96	2.62
CARING - P3D (ours)	13.73	19.92	0.33	0.42	2.12	2.93
CARING - I3D (ours)	18.34	23.60	0.28	0.38	1.55	2.17
Drive&Act - All Classes						
P3D [51]	17.89	21.09	0.28	0.34	1.77	2.12
I3D [52]	11.72	15.97	0.14	0.19	1.10	1.56
P3D + Temperature Scaling [9]	5.89	6.41	0.26	0.30	1.35	1.63
I3D + Temperature Scaling [9]	6.59	8.55	0.12	0.17	0.68	0.99
CARING - P3D (ours)	4.58	5.26	0.24	0.30	1.26	1.57
CARING - I3D (ours)	3.03	6.02	0.11	0.17	0.58	0.9

Standard activity recognition models

Uncertainty-aware models

TABLE I: Reliability of confidence values on the Drive&Act [\[1\]](#) for standard activity recognition models and their extensions with uncertainty-aware calibration algorithms.

true misclassification probability. Additionally, results under the metric of Negative Log Likelihood (NLL) are presented, as high NLL values are strongly correlated to model miscalibration [\[9\]](#). Overall, we average results over the three splits.

B. Confidence estimates for driver activity recognition

In Table I, we compare the described CNN-based activity recognition approaches among themselves and with their uncertainty-aware versions in terms of ECE, Brier Score, and NLL for *overrepresented*, *rare*, and *all* classes in Drive&Act. First, the conjecture has been confirmed that native activity recognition architectures yield unreliable confidence predictions. Evidently, confidence scores produced by I3D score have a misalignment of 15.97% for Drive&Act. Similar problems have been observed in P3D: 21.2% ECE on Drive&Act, an exceedingly high error for safety-critical driving applications.

Model reliability has been clearly improved by learning to attain proper probability estimates, since all of the uncertainty-aware versions exceed the original Softmax values. To our surprise, while I3D has better initial uncertainty estimates than P3D, i.e., 15.97% for I3D in contrast with ECE of 21.09% for P3D, it seems that P3D has a stronger response to both, temperature scaling and CARING approaches than I3D, resulting ECE of 5.26% for CARING-P3D and 6.02% for CARING-I3D. However, this difference is less than 1%, and thereby one would advocate to use I3D due to its higher accuracy in most cases [\[1\]](#), [\[51\]](#), [\[52\]](#). While the expected calibration error is deemed to be essentially important for real applications, one must realize that this metric is complementary to model accuracy and both measures should be taken into consideration when selecting the suitable model for deployment. We further mark, that both temperature scaling and the CARING method *do not influence the model accuracy*, as illustrated in Sections [III-D](#) and [III-E](#). Regarding Pseudo 3D ResNet, we have achieved an overall accuracy of 54.86% on the validation split and 46.62% on the test split of Drive&Act, which does not change after our uncertainty-inspired adaptations. In line

Activity	Number of Samples	Recall (%)	I3D S			CARING-I3D U		
			Mean Conf. (%)	ΔAcc (%)	ECE (%)	Mean Conf. (%)	ΔAcc (%)	ECE (%)
Five most common activities								
sitting_still	2797	95.1	97.96	2.86	2.86	93.84	-1.26	1.84
eating	877	86.42	93.26	6.84	9.33	80.99	-5.43	5.75
fetching_an_object	756	76.03	93.77	17.74	18.28	79.42	3.4	5.32
placing_an_object	688	66.77	93.03	26.25	26.25	75.9	9.13	9.25
reading_magazine	661	92.93	98.58	5.65	6.09	93.35	0.42	2.87
Five most underrepresented activities								
closing_door_inside	30	92.31	98.51	6.21	8.22	86.00	-6.31	8.30
closing_door_outside	22	81.82	93.55	11.73	20.97	86.86	5.04	19.81
opening_backpack	27	0	98.82	98.82	98.82	82.69	82.69	82.69
putting_laptop_into_backpack	26	16.67	92.67	76.00	76.00	76.46	59.8	59.80
taking_laptop_from_backpack	19	0.00	85.25	85.25	85.25	70.08	70.08	70.08

(S) Standard activity recognition models

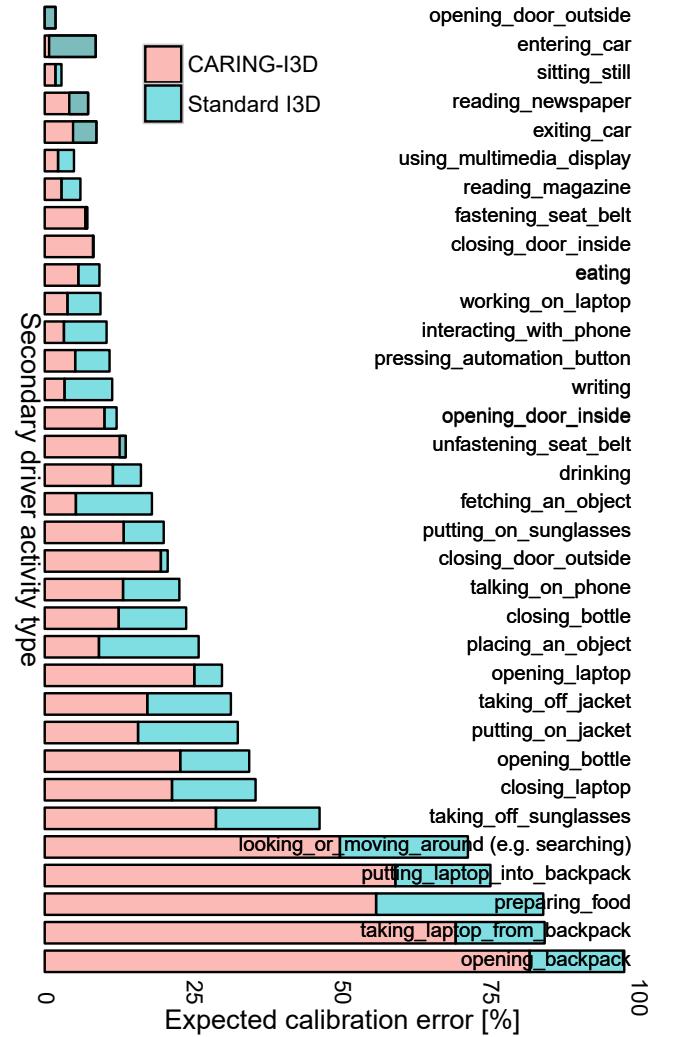
(U) Uncertainty-aware models

TABLE II: Analysis of the resulting confidence estimates of the initial I3D model and its CARING version for individual common and rare Drive&Act activities. *Recall* denotes the recognition accuracy of the current class, while *Mean Conf.* denotes the average confidence estimate produced by the model. Supplemental to the Expected Calibration Error (*ECE*), we report the difference between the mean confidence value and model accuracy (denoted ΔAcc). While in a perfectly calibrated model ΔAcc is 0, ECE is a better evaluation metric, as e.g. if a lot of samples have too high and too low confidence values, their average might lead to a misconception of good calibration. While there is room for improvement for underrepresented and poorly recognized activity classes, the CARING model leads to better uncertainty estimates.

with [1], I3D reaches a higher accuracy of 68.71% on the validation set and 63.09% on the test set.²

It arrives as expected, that the model confidence reliability is correlated to the amount of training data, as seen in the distinguished areas for *common*, *rare*, and *all* classes of Drive&Act (Table I). A case in point lies in the *common classes* setting, which encounters the lowest expected calibration error for both original and uncertainty-aware architectures, i.e., 13.38% for I3D and 5.26% for CARING-I3D. Using intermediate input representation through our CARING calibration network results in the best probability estimates on both datasets and in all evaluation settings. More precisely, with the proposed CARING strategy, our model clearly stands out in front of the raw neural network and the temperature scaling-based one, surpassing them by 9.95% and 2.53% on Drive&Act. The differences are smaller when considering the Brier Score metric, presumably because it is especially insensitive to rare categories [71], which, however, might be of high importance in driver observation, as uncommon events are often linked to distraction. Still, also in terms of the Brier score, the advantages of the CARING model are evident, while the largest performance gain is achieved for underrepresented classes. Overall, our experiments highlight the benefits of learning to attain probability scores *depending on the input*.

Moreover, we assess model performance for the individual classes by examining the top five common and the five least frequent activities in Drive&Act, whose results are displayed in Table II in a separate manner. In addition to ECE, we include the accuracy for samples belonging to the individual classes, the average confidence value achieved with the corresponding model, and the difference between them (denoted as ΔAcc). While a such global confidence-accuracy disagreement



²The slight deviation from the accuracy reported in the original work [1] (between 0.18% and 1.3%) is due to random factors in the training process.

Fig. 6: Exected Calibration Error (ECE) for all 34 secondary activity categories of the Drive&Act dataset (validation split).

is interesting to consider (and is 0 for a perfectly calibrated model), it should be interpreted with care, as it may bring about an imprecise illusion of good confidence calibration, as e.g., many samples with too high and too low confidence values may cancel each other out due to the averaging effects.

Reliability of the confidence scores has been significantly improved via the CARING method and is linked to the amount of training samples and the accuracy. Models have clear issues when learning with few examples, and have a limited performance, e.g., 76.00% for I3D and 59.80% for CARING-I3D ECE in dealing with *putting laptop into backpack*. In contrast, for both, over- and underrepresented classes, the ECE of easy-to-identify activities (i.e., the ones with high accuracy) is much lower. Before calibrated, the average confidence value is always higher than the accuracy (positive ΔAcc), revealing that the models are too optimistic in their predictions. Interestingly, after the CARING conversion is executed, the average model confidence becomes lower than the accuracy for some classes, such as *eating*. In this sense, CARING models are more conservative in their assessment of prediction certainty.

C. Results for different driver behaviours

We now look at different driver behaviour types individually, with the exact ECE for the original classifier and our CARING model provided in Figure 6. We observe in Figure 6, that nearly all secondary behaviours exhibit better quality of the confidence estimates with the proposed CARING model, but certain behaviours still have a large room for improvement.

As previously validated in Table II, the realism of confidences is especially good for more discriminative behaviors (which also obtain higher accuracies) and classes with sufficient training examples. For instance, the activity *eating* has around 800 examples in Drive&Act [1] and a much lower ECE compared to a similar category *preparing food*, of which there are only ~ 50 examples. Especially a mixture of low distinctiveness (e.g., *opening backpack* vs. *taking laptop from backpack*) and low number of training examples are fatal for a category: *taking laptop backpack* is off by 70% in terms of ECE (with CARING), which goes up to 85% without additional confidence calibration. In other words, the model did not correctly predict the class in nearly all cases of this behaviour, while the estimated confidence was very high.

D. Observation results from alternate views and modalities

All activities in Drive&Act are captured by a set of six cameras, five lightweight near-infrared (NIR) cameras and one Kinect camera. The NIR cameras covers several different views from the interior of the vehicle, namely: *front top*, *right top*, *back*, *face view*, and *left top*. In addition, the Kinect camera provides three different modalities: near-infrared, color, and depth. All cameras are synchronized at the frame level, therefore every labeled activity is captured by every camera and corresponds to the same time interval. The principal modality, used in the other evaluation tasks, corresponds to the *front top* view, as done in [1].

In Table III, we show the performance of CARING in all available views and modalities. One of the key differences

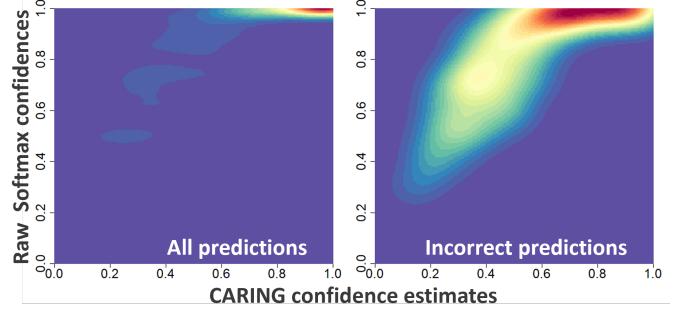


Fig. 7: Distribution of predicted confidences as a 2D histogram for correct and incorrect predictions. The Y axis represents standard Softmax probability estimates, while the X axis are confidences improved through the CARING network. Red denotes common cases (high frequency), while blue illustrates unlikely cases.

between them is the difficulty of the task, e.g., it is easier to recognize an activity in views where the entire body is visible than in views where only the face is visible. We can observe how the *face view*, with its limited aperture, obtains a mean accuracy amongst all classes of only 49.57%, whereas the *front top* obtains an accuracy of 68.71%. Therefore, we can compare how CARING behaves when recognizing the same activities from different levels of difficulty. The ECE of the raw CNN hovers between 10% and 20% on most of the modalities, albeit it shows significant variance and rises up to 80.45% on the *Kinect IR* modality when only common cases are considered. Conversely, the ECE of the CARING models not only is lower, but it shows a very stable behavior with little variance between different views and modalities. Compared to the raw 3D CNN and temperature scaling approach on the main sensor modality of Drive&Act [1], i.e., Front top NIR, CARING improves the validation and test sets ECE by 13.14% and 19.72% for rare classes, 7.65% and 7.16% for common classes and 8.68% and 9.96% for all classes, indicating clear advantages of the proposed CARING model for uncertainty estimation, which is consistent for different amounts of training data (although the overall confidence quality declines with less training examples for all considered approaches). Similar results are observed for other Drive&Act modalities and sensor placements. For example, for the Kinect Depth modality, CARING outperforms the other two methods on validation and test sets by 12.12% and 16.28% for rare classes, 7.07% and 7.10% for common classes and 5.91% and 9.02% for all classes, providing empirical evidence for cross-modal generalization abilities of CARING. Even so, we observe that the ECE of the most challenging view is larger than the ECE of less challenging views, which is consistent with a confidence model where the network is more self-confident on easier tasks.

E. Distribution analysis

We now take a closer look at how the obtained probability estimates are distributed. Figure 7 plots these distributions, where the color highlights the frequency of *Drive&Act* validation set samples obtaining a certain confidence score for

Model	Rare classes						Common classes						All classes											
	ECE (%) val	ECE (%) test	Brier val	Brier test	NLL val	NLL test	Acc (%) val	Acc (%) test	ECE (%) val	ECE (%) test	Brier val	Brier test	NLL val	NLL test	Acc (%) val	Acc (%) test	Brier val	Brier test	NLL val	NLL test	Acc (%) val	Acc (%) test		
Front top NIR																								
Raw 3D CNN	31.48	43.32	0.34	0.47	3.41	4.54	57.02	48.68	10.22	13.38	0.12	0.16	0.90	1.27	80.41	77.5	11.72	15.98	0.14	0.19	0.90	1.56		
Temp. Scaling	24.97	32.38	0.30	0.41	1.96	2.62	55.70	49.78	5.31	6.99	0.11	0.15	0.57	0.83	3.04	3.04	6.59	8.55	0.12	0.17	0.68	0.99		
CARING	18.34	23.60	0.28	0.38	1.55	2.17	2.57		6.22	0.10	0.15	0.50	0.78		3.04	6.02	0.11	0.17	0.58	0.90	68.71	63.09		
Right top NIR																								
Raw 3D CNN	35.89	42.86	0.40	0.48	3.79	4.78	55.57	49.78	14.73	15.83	0.18	0.19	1.28	1.42	72.31	70.06	16.26	18.23	0.19	0.22	1.48	1.73		
Temp. Scaling	25.69	29.52	0.35	0.41	2.20	2.73	55.57	49.78	6.59	6.57	0.15	0.17	0.82	0.93	3.67	3.67	7.71	7.97	0.17	0.20	0.93	1.10		
CARING	20.49	24.03	0.33	0.40	1.92	2.39	3.39		5.30	0.15	0.17	0.76	0.90		3.17	5.67	0.16	0.19	0.85	1.03	63.94	59.92		
Back NIR																								
Raw 3D CNN	40.17	49.47	0.46	0.54	4.42	5.23	44.9	39.95	16.07	19.55	0.22	0.23	1.58	1.67	65.18	66.89	17.85	22.14	0.24	0.26	4.42	2.00		
Temp. Scaling	25.45	33.69	0.39	0.46	2.70	3.00	44.9	39.95	5.19	9.00	0.20	0.20	1.14	1.06	3.17	3.17	6.03	11.00	0.21	0.22	1.26	1.24		
CARING	19.95	26.58	0.37	0.43	2.41	2.67	3.48		5.25	0.19	0.19	1.07	1.01		3.17	6.59	0.20	0.21	1.17	1.16	55.04	53.42		
Face View NIR																								
Raw 3D CNN	39.22	45.56	0.47	0.54	4.60	4.97	42.09	34.92	19.89	21.40	0.29	0.31	2.02	2.02	57.05	49.73	21.34	23.50	0.30	0.33	2.22	2.28		
Temp. Scaling	24.76	26.14	0.40	0.45	2.59	3.03	42.09	34.92	6.30	6.68	0.26	0.28	1.39	1.53	4.63	4.63	6.78	6.81	0.27	0.30	1.48	1.67		
CARING	18.93	23.20	0.38	0.44	2.37	2.94	4.22		0.25	0.28	1.27	1.53			4.63	9.87	0.26	0.30	1.36	1.66	49.57	42.32		
Left top NIR																								
Raw 3D CNN	31.17	44.11	0.34	0.48	2.96	4.23	52.64	47.93	9.43	12.43	0.11	0.16	0.92	1.18	80.74	75.80	11.10	15.16	0.13	0.19	1.08	1.46		
Temp. Scaling	21.46	31.12	0.30	0.41	1.79	2.58	52.64	47.93	4.96	6.23	0.10	0.15	0.59	0.83	2.95	2.95	5.88	7.69	0.12	0.17	0.69	0.99		
CARING	18.14	25.46	0.28	0.40	1.53	2.34	2.45		6.58	0.10	0.15	0.49	0.82			5.77	0.11	0.17	0.57	0.97	66.69	61.87		
Kinect IR																								
Raw 3D CNN	29.65	36.12	0.29	0.45	3.03	3.69	63.83	53.16	80.45	13.69	0.12	0.15	0.84	1.253	82.10	76.94	9.51	15.67	0.13	0.18	1.01	1.48		
Temp. Scaling	20.81	23.88	0.26	0.39	2.00	2.38	63.83	53.16	49.95	6.17	0.11	0.14	0.64	0.89	2.69	2.69	5.11	7.49	0.12	0.16	0.74	1.03		
CARING	15.66	19.90	0.25	0.37	1.78	2.20	2.81		5.19	0.10	0.14	0.53	0.85			6.16	0.11	0.16	0.63	0.97	72.96	65.05		
Kinect Depth																								
Raw 3D CNN	32.43	41.69	0.34	0.47	3.23	4.44	61.34	48.80	11.72	16.83	0.15	0.21	1.01	1.51	77.47	70.86	13.13	18.94	0.16	0.23	1.17	1.78		
Temp. Scaling	23.88	29.36	0.30	0.41	1.94	2.66	61.34	48.80	4.93	6.81	0.14	0.19	0.70	1.03	3.44	3.44	5.82	8.32	0.15	0.21	0.80	1.18		
CARING	18.84	25.05	0.28	0.39	1.67	2.35	3.08		5.00	0.13	0.18	0.65	0.97			5.50	0.14	0.20	0.72	1.10	69.40	59.83		
Kinect Color																								
Raw 3D CNN	28.00	40.38	0.35	0.42	3.14	4.66	56.42	53.40	9.80	13.21	0.12	0.18	0.98	1.20	80.27	72.00	8.99	15.56	0.14	0.20	1.14	1.51		
Temp. Scaling	21.47	31.06	0.31	0.37	1.80	2.66	56.42	53.40	4.93	7.64	0.12	0.16	0.60	0.78	3.08	3.08	5.86	9.43	0.13	0.18	0.69	0.95		
CARING	15.88	24.10	0.30	0.35	1.80	2.21	2.73		2.73	6.10	0.11	0.16	0.51	0.74			3.08	6.54	0.12	0.18	0.59	0.87	68.34	62.70

TABLE III: Impact of the confidence adjustment through temperature scaling and our CARING model for different modalities and views inside the vehicle cabin. *Raw 3D CNN* indicates I3D without temperature scaling, *Temp. Scaling* indicates I3D equipped with the method of [9] and *CARING* is our input-guided framework. Additionally to the confidence quality metrics (ECE, Brier Score and NLL, where ECE is considered the most reliable measure), we report the classification accuracies of the individual modalities for the I3D network on the Drive&Act dataset. We evaluate our model not only on all classes but also on rare classes (underrepresented) and common classes (overrepresented). Keep in mind, that as both models do not impact the order of the logits, the *accuracy of the model does not change*. In the vast majority of modalities, our confidence improvement through CARING leads to significantly more realistic confidence estimates, indicating its effectiveness across different sensor types and camera placements inside the vehicle cabin.

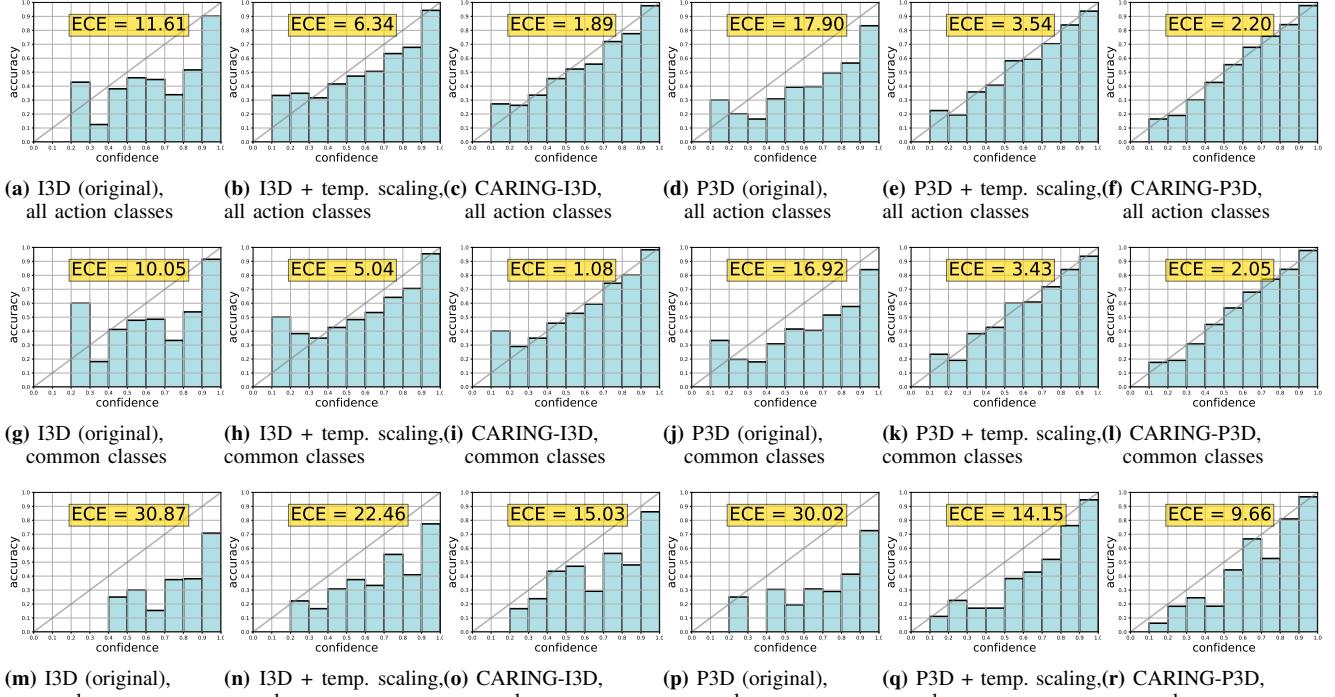


Fig. 8: Reliability diagrams of different models reflect the agreement between the confidence values and the empirically measured probability of correct prediction (results of one Drive&Act validation split). A model with *perfectly calibrated uncertainty scores would match the diagonal* (a detailed explanation in Section IV-E). Note, that the ECE values deviate from Table I, as they visualize a single split, while the final reported results are averaged over all splits. While the temperature scaling consistently improves the confidence estimates, our CARING model leads to the lowest calibration error in all settings.

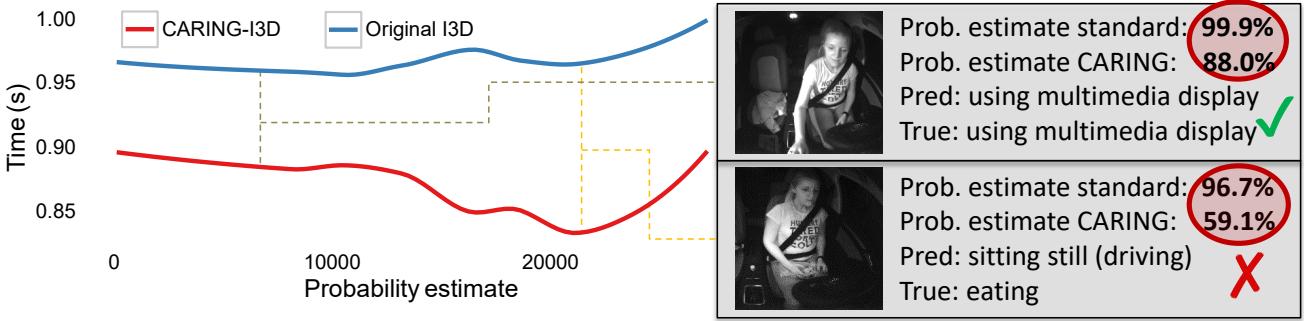


Fig. 9: Example of average confidence estimates (median value with a sliding window range of one minute) during a single driving session for the original (blue line) and CARING (red line) models and qualitative examples of a correct and an incorrect outcome (on the right).

correct and incorrect predictions. The X axis denotes the improved probability estimates after the CARING transformation, while the Y axis corresponds the original *Softmax* confidences. For correct predictions (Figure 7 on the left), model confidences tend to be high both with and without CARING, but the CARING confidences have much higher variance, which is visible through the horizontal spread of the density plot. This effect is even stronger in case the model made a mistake (Figure 7 on the right): the red high-frequency region is much wider on the X-axis than on the Y-axis, highlighting more reserved confidence values for CARING.

Figure 8 visualizes the agreement between the estimated model confidence and the empirically measured probability of the correct outcome by using *reliability diagrams* (explained in Section III-A). In case of good estimates, the result will be closely approaching the diagonal line. Values above the diagonal are correlated to models being over-confident in their prediction, while values below imply that the model harbours suspicions about the outcome and the accurate prediction probability is higher than assumed.

We first analyze the reliability diagrams of the original action recognition networks. Both P3D and I3D confidence values largely deviate from the target, showing clear biases towards excessively optimistic scores (i.e., values are oftentimes below the diagonal in Figures 8a, 8d, 8g, 8j, 8m, 8p). One exception worth discussion is an above-diagonal peak in the low probability segment for *all* and *common* classes, hinting that in “easier” settings, low confidence samples often turn out to be correct (8a, 8d, 8g, 8j). Still, in the “harder” setting of *rare* activities (Figure 8m, 8p), the biases towards too high probabilities are noticeable for all values.

We have observed a clear positive effect of temperature scaling (Figures 8b, 8e, 8h, 8k, 8n, 8q) and our CARING model (Figures 8c, 8f, 8i, 8l, 8o, 8r). Evidently, CARING models outperform other approaches in all settings and lead to almost perfect reliability diagrams for *all* and *common* classes. Yet, both temperature scaling and CARING strategies still have issues in dealing with *rare* classes, with confidence being too high, indicating a vital direction for future explorations.

It is worth noting that ECE might be in a slight disarray with the visual reliability diagram representation, as the metric weighs the misalignment in each bin by the amount of data points in it, whereas the reliability diagrams do not manifest such frequency distribution. For instance, while the CARING-

I3D model in Figure 8i marginally exceeds the target diagonal, it has lower expected calibration error than CARING-P3D which seems to yield nearly perfect results in Figure 8l. As there are only very few examples in the low-confidence bin, they are overshadowed by smaller differences in the high-confidence bins, which therefore contribute more to the final.

V. DISCUSSION: TOWARDS RELIABLE AND SECURE DRIVER OBSERVATION

A. Significance for Driver Observation

Although CARING consistently improves the realism of confidence values, the results vary greatly depending on the behaviour type (Figure 6). As it comes to real-life applications inside the vehicle, the maturity level of the proposed model depends on the concrete use-case. Drive&Act is a fine-grained dataset with both, manual and fully automated driving, and therefore covers a high number of secondary activities. Some of these behaviours were almost unthinkable behind the steering wheel until now, but will become more common in the future (such as *working on laptop*). For the use-case of identifying driver distraction, the most significant Drive&Act category is *sitting still*, which comprises both, (1) driving and (2) sitting with the hands not on the steering wheel, while the car is in the self-driving mode (without any distractive activities). This category is inherently linked to lower distraction and distinguishing it from other driver states is crucial for identifying distraction at SAE levels 0 to 3. Recognition of this behaviour, as well as most of other behaviours prevalent in manual driving, (e.g., *using multimedia display*, *fastening seat belt*) consistently yields high-quality confidence estimates. The rather long-term use-case of increased comfort through automatic driver-centered adaptation, (e.g., softer driving if the person is drinking tea or sleeping at SAE levels 4 and 5), on the other hand, might require finer granularity of driver behaviours. In this case, the readiness level of our system strongly depends on the concrete application and significant improvement would still be required if very concise recognition is desired (for example, *opening vs. closing bottle*). Nevertheless, the most fundamental categories, such as *eating*, *reading* or *interacting with phone* achieve $ECE < 10\%$, which might be sufficient for most of the applications.

Our experiments indicate that previously used driver observation models do not provide proper uncertainty estimates, therefore ignoring their own limits. Figure 9 illustrates an example of average confidence estimates during a single driving session of the Drive&Act dataset for the original (blue line) and CARING (red line) models. The original model consistently stays in the area $> 95\%$, indicating strong overconfidence, since the average Drive&Act accuracy is clearly below this value. Two qualitative examples on the right showcase very high confidence values (99.9% and 96.7%) of the original model in cases of both, correct predictions and misclassifications. Such overly confident models provide a false sense of security and are dangerous if affecting the overall driving dynamics. Furthermore, such false-positives predictions are very damaging for the driving experience and human-vehicle interaction. In the vast majority of cases, the uncertainty error was towards overconfidence, which is clearly improved though our CARING model. In the example of 9, CARING clearly lowers the average confidence estimate (red line). In the qualitative examples on the right, CARING estimate is much lower in case of misclassification (59.1% vs. 88.0%). Still, one might argue that a value of 59.1% is relatively high and finding a good threshold for acceptance of the recognition result remains an open question.

Classification uncertainty plays a significant role in driver observation, as the the models are deployed in a dynamic world where domain shifts, novel behaviours and other unforeseen situations may occur at any time. Besides being potentially dangerous, false-positives caused by overly confident models are often highly disturbing for the user. For example, in the use-case of driver-centered adaptation, repeated false recognition that the human is, *e.g.*, reading triggering a reaction like turning on the light would strongly bother the user. Apart from the direct benefits of proper confidence values for decision-making systems, good assessment of uncertainty enhances model interpretability. For example, in the realistic scenario of open-world recognition, low-confidence input might be passed to human experts, which would provide the correct annotations (*i.e.* active learning) and therefore improve the decision boundary. In summary, obtaining realistic uncertainty estimates is important in real-life driving systems due to (1) safety, (2) better human-machine interaction (through avoidance of false positives triggering an unwanted response), and (3) better interpretability.

B. Discussion of Limitations

While our method leads to much more reliable confidence estimates in driver behaviour analysis, it is not without limitations. First, we observe, that the accuracy-confidence mismatch is still relatively large in case of activities with little training examples. Secondly, while we achieve the best results with the proposed CARING approach, we acknowledge, that this comes with increased computational cost due to an additional external calibration network. This overhead of course strongly depends on the architecture choice and is comparably small in our case (66.7K Parameters and 0.13 MFLOPS for our two-layer calibration network using 1024 dimensional

input representation vector holding the intermediate I3D representations). Furthermore, *identifying misclassifications* among the training classes is not the same as *identifying novelty* [78]. The essence of calibration-based methods, such as temperature scaling or our CARING model, is learning realistic confidence values on a held-out validation set. As shown in a recent study from the area of image classification [78], the reliance on this held-out validation set also becomes the greatest weakness of such models when facing domain shifts. In other words, confidence calibration is effective, as long as the test data roughly reflects the distribution of the validation set. Producing probability estimates of the observed behavior class that was previously unseen during training is therefore a different challenge that we aim to address in the future. Nevertheless, our work takes a step towards driver observation models which are not only accurate, but also to determine, how likely they are to be right in their prediction, which is vital for reliable and transparent decision systems within the vehicle.

C. Outlook: Reliable and Secure Driver Observation

The goal of driver observation is to delegate more actions and decisions to the vehicle itself and reacting to the captured driver states accordingly. As in any decision process, the reliability and security of the incoming information become important. Our current work provides an accurate method for measuring the uncertainty of this information. Taking this as a starting point, one could follow different research directions which are currently rather overlooked.

Adversarial attacks are one of such examples. Though thoroughly explored for signals coming from the exterior of the vehicle [82], [83], it is still not known how extensive the risk is of an attack against the cameras inside the vehicle cabin. Unfortunately, deep learning research suggests that raw Softmax confidences are highly unreliable in case of adversarial attacks [12]. Analyzing the uncertainty of the networks can reduce the effect of such attacks and provide a way to detect them. Still, how much additional effort is indeed needed to penetrate this second layer of protection for action recognition models remains an open question.

A related direction is that of *interpretability* [84]. As we move towards end-to-end approaches, it is easier for dataset biases to influence predictions. In the case of driver observation, a movement of the hand towards the mouth could be for example misinterpreted as eating [60]. Visualizing the features which determine the predictions can help build trust and ensure that situations as the aforementioned will not occur. The CARING model allows for a second interpretation, namely that of the uncertainty estimates which could lead to additional insights into model reasoning.

Last but not least comes *security and privacy* [85], [86]. Monitoring the inside of a vehicle exposes a lot of private sensitive information. Aside form the video feed itself, the processed information is also of high concern, especially when used to evaluate professional drivers, resolve accident disputes, or used for marketing purposes. In all those cases, knowing the reliability of predictions can help to better quantify the benefits and dangers of using the driver observation information.

VI. CONCLUSION

While conventional driver observation research targets high-accuracy frameworks in the first place, our goal is to develop driver activity recognition models which are not only accurate, but also produce uncertainty estimates that can safely be interpreted as probabilities of a correct outcome. First, we study the reliability of model confidences of existing driver activity recognition models and come to worrying conclusions: there is a significant discrepancy between the original confidence estimates of the networks and the empirically measured probabilities of a correct prediction. To remedy this, we implement two methods for calibrating the network estimates: leveraging the temperature scaling approach often used in image recognition for such purposes and our newly proposed model, which we refer to as CARING. Our CARING model learns to dynamically generate values used for scaling the model confidences based on input representation with an external calibration network. Our extensive experiments indicate strong improvement through temperature scaling alone, while our CARING model consistently yields the best results. Our study provides encouraging evidence, that the proposed uncertainty-aware models can be efficiently used to develop reliable driver observation systems.

REFERENCES

- [1] M. Martin, A. Roitberg, M. Haurilet, M. Horne, S. Reiß, M. Voit, and R. Stiefelhagen, “Drive&act: A multi-modal dataset for fine-grained driver behavior recognition in autonomous vehicles,” in *Proc. ICCV*, 2019, pp. 2801–2810.
- [2] E. Ohn-Bar and M. M. Trivedi, “Hand gesture recognition in real time for automotive interfaces: A multimodal vision-based approach and evaluations,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 15, no. 6, pp. 2368–2377, 2014.
- [3] M. Martin, J. Popp, M. Anneken, M. Voit, and R. Stiefelhagen, “Body pose and context information for driver secondary task detection,” in *Proc. IV*, 2018, pp. 2015–2021.
- [4] A. Jain, H. S. Koppula, B. Raghavan, S. Soh, and A. Saxena, “Car that knows before you do: Anticipating maneuvers via learning temporal driving models,” in *Proc. ICCV*, 2015, pp. 3182–3190.
- [5] Z. Wharton, A. Behera, Y. Liu, and N. Bessis, “Coarse temporal attention network (CTA-Net) for driver’s activity recognition,” in *Proc. WACV*, 2021, pp. 1279–1289.
- [6] P. Weyers, D. Schiebener, and A. Kummert, “Action and object interaction recognition for driver activity classification,” in *Proc. ITSC*, 2019, pp. 4336–4341.
- [7] P. Li, M. Lu, Z. Zhang, D. Shan, and Y. Yang, “A novel spatial-temporal graph for skeleton-based driver action recognition,” in *Proc. ITSC*, 2019, pp. 3243–3248.
- [8] P. Gebert, A. Roitberg, M. Haurilet, and R. Stiefelhagen, “End-to-end prediction of driver intention using 3D convolutional neural networks,” in *Proc. IV*, 2019, pp. 969–974.
- [9] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, “On calibration of modern neural networks,” in *Proc. ICML*, 2017, pp. 1321–1330.
- [10] N. Sünderhauf, O. Brock, W. J. Scheirer, R. Hadsell, D. Fox, J. Leitner, B. Upcroft, P. Abbeel, W. Burgard, M. Milford, and P. Corke, “The limits and potentials of deep learning for robotics,” *The International Journal of Robotics Research*, vol. 37, no. 4–5, pp. 405–420, 2018.
- [11] D. Hendrycks and K. Gimpel, “A baseline for detecting misclassified and out-of-distribution examples in neural networks,” in *Proc. ICLR*, 2017.
- [12] A. Nguyen, J. Yosinski, and J. Clune, “Deep neural networks are easily fooled: High confidence predictions for unrecognizable images,” in *Proc. CVPR*, 2015, pp. 427–436.
- [13] J. D. Ortega, N. Kose, P. Cañas, M.-A. Chao, A. Unnervik, M. Nieto, O. Otaegui, and L. Salgado, “DMD: A large-scale multi-modal driver monitoring dataset for attention and alertness analysis,” in *Proc. ECCV*, 2020, pp. 387–405.
- [14] L. Fontanari, M. Gonzalez, G. Vallortigara, and V. Girotto, “Probabilistic cognition in two indigenous mayan groups,” *Proceedings of the National Academy of Sciences*, vol. 111, no. 48, pp. 17 075–17 080, 2014.
- [15] Y. Gal and Z. Ghahramani, “Dropout as a bayesian approximation: Representing model uncertainty in deep learning,” in *Proc. ICML*, 2016, pp. 1050–1059.
- [16] *Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles SAE J 3016*, Society of Automotive Engineers (SAE) Std., 2016.
- [17] J. Radlmayr, C. Gold, L. Lorenz, M. Farid, and K. Bengler, “How traffic situations and non-driving related tasks affect the take-over quality in highly automated driving,” in *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 58, no. 1, 2014, pp. 2063–2067.
- [18] J. Ludwig, M. Martin, M. Horne, M. Flad, M. Voit, R. Stiefelhagen, and S. Hohmann, “Driver observation and shared vehicle control: supporting the driver on the way back into the control loop,” in *Automatisierungstechnik*, vol. 66(2), pp. 146 – 159, 2018.
- [19] S. International, “Automated driving: Levels of driving automation are defined in new SAE international standard J3016,” 2014.
- [20] M. Flad, P. Karg, A. Roitberg, M. Martin, M. Mazewitsch, C. Lange, E. Kenar, L. Ahrens, B. Flecken, L. Kalb, B. Karakaya, J. Ludwig, A. Prusich, R. Stiefelhagen, and S. Hohmann, “Personalisation and control transition between automation and driver in highly automated cars,” in *Smart Automotive Mobility*, 2020, pp. 1–70.
- [21] E. Ohn-Bar and M. M. Trivedi, “Looking at humans in the age of self-driving and highly automated vehicles,” *IEEE Transactions on Intelligent Vehicles*, vol. 1, no. 1, pp. 90–104, 2016.
- [22] E. Wolf, M. Martinez, A. Roitberg, R. Stiefelhagen, and B. Deml, “Estimating mental load in passive and active tasks from pupil and gaze changes using bayesian surprise,” in *Proc. ICMI-MCPMD*, 2018, p. 6.
- [23] N. Deo and M. M. Trivedi, “Looking at the driver rider in autonomous vehicles to predict take-over readiness,” *IEEE Transactions on Intelligent Vehicles*, vol. 5, no. 1, pp. 41–52, 2019.
- [24] B. Mok, M. Johns, K. J. Lee, D. Miller, D. Sirkin, P. Ive, and W. Ju, “Emergency, automation off: Unstructured transition timing for distracted drivers of automated vehicles,” in *Proc. ITSC*, 2015, pp. 2458–2464.
- [25] B. Gleeson, K. MacLean, A. Haddadi, E. Croft, and J. Alcazar, “Gestures for industry intuitive human-robot communication from human observation,” in *Proc. HRI*, 2013, pp. 349–356.
- [26] A. Roitberg, N. Soman, A. Perzylo, M. Rickert, and A. Knoll, “Multi-modal human activity recognition for industrial manufacturing processes in robotic workcells,” in *Proc. ICMI*, 2015, pp. 259–266.
- [27] S. Singh, “Critical reasons for crashes investigated in the national motor vehicle crash causation survey,” Stats (National Highway Traffic Safety Administration, Washington, DC), Report No. DOT HS 812 115., Tech. Rep., 2015.
- [28] D. Lloyd, D. Wilson, D. Mais, W. Deda, and A. Bhagat, “Reported road casualties great britain: 2014 annual report,” 2015.
- [29] J. Fang, D. Yan, J. Qiao, J. Xue, and H. Yu, “DADA: Driver attention prediction in driving accident scenarios,” *IEEE Transactions on Intelligent Transportation Systems*, 2021.
- [30] Y. Gal, “Uncertainty in deep learning,” *University of Cambridge*, 2016.
- [31] A. Malinin, “Uncertainty estimation in deep learning with application to spoken language assessment,” Ph.D. dissertation, University of Cambridge, 2019.
- [32] J. Nixon, M. W. Dusenberry, L. Zhang, G. Jerfel, and D. Tran, “Measuring calibration in deep learning,” in *CVPRW*, 2019, pp. 38–41.
- [33] A. P. Dawid, “The well-calibrated bayesian,” *Journal of the American Statistical Association*, vol. 77, no. 379, pp. 605–610, 1982.
- [34] A. H. Murphy and E. S. Epstein, “Verification of probabilistic predictions: A brief review,” *Journal of Applied Meteorology and Climatology*, vol. 6, no. 5, pp. 748–755, 1967.
- [35] A. Roitberg, M. Haurilet, M. Martinez, and R. Stiefelhagen, “Uncertainty-sensitive activity recognition: A reliability benchmark and the CARING models,” in *Proc. ICPR*, 2021, pp. 3814–3821.
- [36] I. Jegham, A. B. Khalifa, I. Alouani, and M. A. Mahjoub, “Soft spatial attention-based multimodal driver action recognition using deep learning,” *IEEE Sensors Journal*, vol. 21, no. 2, pp. 1918–1925, 2021.
- [37] Y. Xing, C. Lv, H. Wang, D. Cao, E. Velenis, and F.-Y. Wang, “Driver activity recognition for intelligent vehicles: A deep learning approach,” *IEEE Transactions on Vehicular Technology*, vol. 68, no. 6, pp. 5379–5390, 2019.
- [38] A. Ragab, C. Craye, M. S. Kamel, and F. Karray, “A visual-based driver distraction recognition and detection using random forest,” in *Proc. ICIAR*, 2014, pp. 256–265.

- [39] N. Das, E. Ohn-Bar, and M. M. Trivedi, "On performance evaluation of driver hand detection algorithms: Challenges, dataset, and metrics," in *Proc. ITSC*, 2015, pp. 2953–2958.
- [40] M. J. Rahman, S. S. Beauchemin, and M. A. Bauer, "Predicting driver behaviour at intersections based on driver gaze and traffic light recognition," *IET Intelligent Transport Systems*, vol. 14, no. 14, pp. 2083–2091, 2021.
- [41] Y. Liang, M. L. Reyes, and J. D. Lee, "Real-time detection of driver cognitive distraction using support vector machines," *IEEE Transactions on Intelligent Transportation Systems*, vol. 8, no. 2, pp. 340–350, 2007.
- [42] E. Murphy-Chutorian and M. M. Trivedi, "Head pose estimation and augmented reality tracking: An integrated system and evaluation for monitoring driver awareness," *IEEE Transactions on Intelligent Transportation Systems*, vol. 11, no. 2, pp. 300–311, 2010.
- [43] E. Murphy-Chutorian, A. Doshi, and M. M. Trivedi, "Head pose estimation for driver assistance systems: A robust algorithm and experimental evaluation," in *Proc. ITSC*, 2007, pp. 709–714.
- [44] C. Tran, A. Doshi, and M. M. Trivedi, "Modeling and prediction of driver behavior by foot gesture analysis," *Computer Vision and Image Understanding*, vol. 116, no. 3, pp. 435–445, 2012.
- [45] L. Bi, Y. Lu, X. Fan, J. Lian, and Y. Liu, "Queuing network modeling of driver eeg signals-based steering control," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 25, no. 8, pp. 1117–1124, 2016.
- [46] E. Ohn-Bar, S. Martin, A. Tawari, and M. M. Trivedi, "Head, eye, and hand patterns for driver activity recognition," in *Proc. ICPR*, 2014, pp. 660–665.
- [47] L. Xu and K. Fujimura, "Real-time driver activity recognition with random forests," in *Proc. AutomotiveUI*, 2014, pp. 1–8.
- [48] M. Martin, M. Voit, and R. Stiefelhagen, "Dynamic interaction graphs for driver activity recognition," in *Proc. ITSC*, 2020, pp. 1–7.
- [49] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. CVPR*, 2016, pp. 770–778.
- [50] X. Wang, R. Girshick, A. Gupta, and K. He, "Non-local neural networks," in *Proc. CVPR*, 2018, pp. 7794–7803.
- [51] Z. Qiu, T. Yao, and T. Mei, "Learning spatio-temporal representation with pseudo-3D residual networks," in *Proc. ICCV*, 2017, pp. 5534–5542.
- [52] J. Carreira and A. Zisserman, "Quo vadis, action recognition? A new model and the kinetics dataset," in *Proc. CVPR*, 2017, pp. 4724–4733.
- [53] D. Feng, C. Haase-Schütz, L. Rosenbaum, H. Hertlein, C. Glaeser, F. Timm, W. Wiesbeck, and K. Dietmayer, "Deep multi-modal object detection and semantic segmentation for autonomous driving: Datasets, methods, and challenges," *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 3, pp. 1341–1360, 2021.
- [54] K. Peng, J. Fei, K. Yang, A. Roitberg, J. Zhang, F. Bieder, P. Heidenreich, C. Stiller, and R. Stiefelhagen, "MASS: Multi-attentional semantic segmentation of LiDAR data for dense top-view understanding," *IEEE Transactions on Intelligent Transportation Systems*, 2022.
- [55] C. Ma, J. Zhang, K. Yang, A. Roitberg, and R. Stiefelhagen, "DenseP-ASS: Dense panoramic semantic segmentation via unsupervised domain adaptation with attention-augmented context exchange," in *Proc. ITSC*, 2021, pp. 2766–2772.
- [56] S. M. Kouchak and A. Gaffar, "Detecting driver behavior using stacked long short term memory network with attention layer," *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 6, pp. 3420–3429, 2021.
- [57] L.-W. Chen and H.-M. Chen, "Driver behavior monitoring and warning with dangerous driving detection based on the internet of vehicles," *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 11, pp. 7232–7241, 2021.
- [58] S. Reiß, A. Roitberg, M. Haurilet, and R. Stiefelhagen, "Activity-aware attributes for zero-shot driver behavior recognition," in *Proc. CVPRW*, 2020, pp. 902–903.
- [59] B. Qin, J. Qian, Y. Xin, B. Liu, and Y. Dong, "Distracted driver detection based on a CNN with decreasing filter size," *IEEE Transactions on Intelligent Transportation Systems*, 2021.
- [60] A. Roitberg, M. Haurilet, S. Reiß, and R. Stiefelhagen, "CNN-based driver activity understanding - Shedding light on deep spatiotemporal representations," in *Proc. ITSC*, 2020, pp. 1–6.
- [61] S. Reiß, A. Roitberg, M. Haurilet, and R. Stiefelhagen, "Deep classification-driven domain adaptation for cross-modal driver behavior recognition," in *Proc. IV*, 2020, pp. 1042–1047.
- [62] Y. Abouelnaga, H. M. Eraqi, and M. N. Moustafa, "Real-time distracted driver posture classification," in *Proc. NeurIPS-MLITS*, 2018.
- [63] O. Kopuklu, N. Kose, A. Gunduz, and G. Rigoll, "Resource efficient 3D convolutional neural networks," in *Proc. ICCVW*, 2019, pp. 1910–1919.
- [64] O. Köpüklü, T. Ledwon, Y. Rong, N. Kose, and G. Rigoll, "DriverMHG: A multi-modal dataset for dynamic recognition of driver micro hand gestures and a real-time recognition framework," in *Proc. FG*, 2020, pp. 77–84.
- [65] O. Olabiyi, E. Martinson, V. Chintalapudi, and R. Guo, "Driver action prediction using deep (bidirectional) recurrent neural network," *arXiv preprint arXiv:1706.02257*, 2017.
- [66] D. Zhou, H. Liu, H. Ma, X. Wang, X. Zhang, and Y. Dong, "Driving behavior prediction considering cognitive prior and driving context," *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 5, pp. 2669–2678, 2021.
- [67] A. Niculescu-Mizil and R. Caruana, "Predicting good probabilities with supervised learning," in *Proc. ICML*, 2005, pp. 625–632.
- [68] M. H. DeGroot and S. E. Fienberg, "The comparison and evaluation of forecasters," *Journal of the Royal Statistical Society: Series D (The Statistician)*, vol. 32, no. 1-2, pp. 12–22, 1983.
- [69] J. Platt, "Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods," *Advances in Large Margin Classifiers*, vol. 10, no. 3, pp. 61–74, 1999.
- [70] A. Bansal, A. Farhadi, and D. Parikh, "Towards transparent systems: Semantic characterization of failure modes," in *Proc. ECCV*, 2014, pp. 366–381.
- [71] B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles," in *Proc. NeurIPS*, 2017, pp. 6402–6413.
- [72] B. Zadrozny and C. Elkan, "Obtaining calibrated probability estimates from decision trees and naive bayesian classifiers," in *Proc. ICML*, 2001, pp. 609–616.
- [73] ———, "Transforming classifier scores into accurate multiclass probability estimates," in *Proc. KDD*, 2002, pp. 694–699.
- [74] M. Ott, M. Auli, D. Grangier, and M. Ranzato, "Analyzing uncertainty in neural machine translation," in *Proc. ICML*, 2018, pp. 3953–3962.
- [75] M. P. Naeini, G. Cooper, and M. Hauskrecht, "Obtaining well calibrated probabilities using bayesian binning," in *Proc. AAAI*, 2015.
- [76] Y. Huang, W. Li, F. Macharet, R. A. Gabriel, and L. Ohno-Machado, "A tutorial on calibration measurements and calibration models for clinical prediction models," *Journal of the American Medical Informatics Association*, vol. 27, no. 4, pp. 621–633, 2020.
- [77] G. W. Brier, "Verification of forecasts expressed in terms of probability," *Monthly Weather Review*, vol. 78, no. 1, pp. 1–3, 1950.
- [78] Y. Ovadia, E. Fertig, J. Ren, Z. Nado, D. Sculley, S. Nowozin, J. Dillon, B. Lakshminarayanan, and J. Snoek, "Can you trust your model's uncertainty? evaluating predictive uncertainty under dataset shift," in *Proc. NeurIPS*, pp. 13 991–14 002.
- [79] J. Kruppa, Y. Liu, H.-C. Diener, T. Holste, C. Weimar, I. R. König, and A. Ziegler, "Probability estimation with machine learning methods for dichotomous and multicategory outcome: Applications," *Biometrical Journal*, vol. 56, no. 4, pp. 564–583, 2014.
- [80] A. Kumar, S. Sarawagi, and U. Jain, "Trainable calibration measures for neural networks from kernel mean embeddings," in *Proc. ICML*, 2018, pp. 2805–2814.
- [81] J. Mukhoti, V. Kulharia, A. Sanyal, S. Golodetz, P. Torr, and P. Dokania, "Calibrating deep neural networks using focal loss," *Advances in Neural Information Processing Systems*, vol. 33, pp. 15 288–15 299, 2020.
- [82] Y. Deng, X. Zheng, T. Zhang, C. Chen, G. Lou, and M. Kim, "An analysis of adversarial attacks and defenses on autonomous driving models," in *Proc. PerCom*, 2020, pp. 1–10.
- [83] Y. Li, X. Xu, J. Xiao, S. Li, and H. Shen, "Adaptive square attack: Fooling autonomous cars with adversarial traffic signs," *IEEE Internet of Things Journal*, vol. 8, pp. 6337–6347, 2021.
- [84] E. Zablocki, H. Ben-Younes, P. Pérez, and M. Cord, "Explainability of vision-based autonomous driving systems: Review and challenges," *arXiv preprint arXiv:2101.05307*, 2021.
- [85] L. Collingwood, "Privacy implications and liability issues of autonomous vehicles," *Information & Communications Technology Law*, vol. 26, no. 1, pp. 32–45, 2017.
- [86] K. Ren, Q. Wang, C. Wang, Z. Qin, and X. Lin, "The security of autonomous driving: Threats, defenses, and future directions," *Proceedings of the IEEE*, vol. 108, no. 2, pp. 357–372, 2020.